arXiv:1510.02132v1 [math.CO] 7 Oct 2015

Discrete Envy-free Division of Necklaces and Maps

Roberto Barrera[1], Kathryn Nyman[2], Amanda Ruiz[3], Francis Edward Su[4] and Yan X Zhang[5]

[1]Department of Mathematics, Texas A&M University
[2]Department of Mathematics, Willamette University
[3]Department of Mathematics, University of San Diego
[4]Department of Mathematics, Harvey Mudd College
[5]Department of Mathematics, UC Berkeley

October 9, 2015

Abstract

We study the discrete variation of the classical cake-cutting problem where n players divide a 1-dimensional cake with exactly $(n-1)$ cuts, replacing the continuous, infinitely divisible "cake" with a necklace of discrete, indivisible "beads." We focus specifically on *envy-free* divisions, exploring different constraints on player-preferences. We show we usually cannot guarantee an *envy-free* division and consider situations where we can obtain an ϵ-envy-free division for ϵ relatively small. We also prove a 2-dimensional result with a grid of indivisible objects. This may be viewed as a way to divide a state with indivisible districts among a set of constituents, producing somewhat gerrymandered regions that form an envy-free division of the state.

1 Introduction

The classical 1-dimensional cake cutting problem is fairly well-studied in the literature. Formally: suppose a rectangular continuous cake is to be divided among n players, where each player may have different notions of what parts of the cake are valuable. The cake is to be divided by $(n-1)$ *cuts* parallel to the left edge of the cake, creating n *pieces* (intervals) of cake. A choice of such a set of cuts (or equivalently, the set of pieces) is called a *division* of the cake. Modeling the cake as the interval $[0, l]$ and letting $x_i \geq 0$ be the length of the i-th piece of a division, the *space of divisions* is given by n-tuples $(x_1, \ldots, x_n), x_i \geq 0$, such that $x_1 + x_2 + \ldots + x_n = l$. Thus the set of divisions of the cake is naturally parametrized by a n-simplex. A point in the simpex corresponds to a division

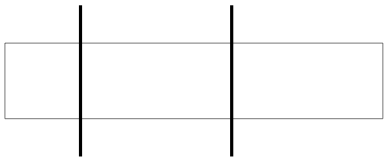

Figure 1: A division of a continuous cake, cut by a cutset of 2 cuts into 3 pieces.

into pieces of width $x_1, x_2, \ldots, x_n$. We call the set of cuts corresponding to such a point a *cutset*. For an example of this setup, see Figure 1.

In our setup, we assume each player i has a *valuation* function v_i defined on any possible piece of cake, having as its codomain the **nonnegative** real numbers. We require the valuation function to *additive*, meaning if a player has valuation a for a piece $[x, y]$ and b for a piece $[y, z]$, then the player values the piece $[x, z]$ at $(a + b)$. This gives rise to a transitive *preference* relation between each two possible pieces of cake by having player i prefer piece A to B if and only if $v_i(A) \geq v_i(B)$. We stress that as a technical term, a player may both prefer A to B and B to A, liking them equally, which is a non-intuitive use of the word. If we assert that A is preferred to B but not vice-versa, we say that the player *strictly prefers* A to B. In the context of a particular division, a player *prefers* a piece of cake if the player prefers the piece to any other piece in the division (note a player may prefer several pieces of cake simultaneously). An *allocation* of the cake is a division equipped with a bijection between the players and the pieces (in other words, an assignment of a different piece to every player).

It is typically assumed that if a player prefers x_{ij} for all j in a sequence of cutsets $(x_{11}, \ldots, x_{n1}), (x_{12}, \ldots, x_{n2}), \ldots$ converging to $(x_1, \ldots, x_n)$, then the player prefers x_i in the limiting cutset. Note this "continuity condition" implies some properties of the valuation function; in particular, it implies that there must be no point mass with positive valuation (so it justifies breaking the closed interval cake into open intervals and calling the result a "division" of the cake).

The archetypal fair-division problem looks at the space of allocations and seeks to prove the existence of and/or find allocations with desirable properties. One common example is calling an allocation *envy-free* if every player prefers the piece they are assigned to by the allocation (i.e. they do not *envy* any other player in the sense that they strictly prefer another player's allocated piece to their own). Under these assumptions the cake cutting problem has a solution [8]:

Theorem 1 (Envy-Free Cake-Cutting)**.** *For any set of n players's preferences, there exists an envy-free allocation of cake using* $(n - 1)$ *cuts. Furthermore, there exists a finite* ϵ*-approximate algorithm.*

The existence of envy-free divisions has been known since Neyman [6], with recent attention paid to finding constructive proofs with potential for applica-

tions. The first constructive n-player envy-free solution is due to Brams and Taylor [3]; it was was finite but unbounded – meaning that it would take a finite number of steps but that number may be arbitrarily large! Moreover, the cake could be cut into an arbitrarily large number of pieces.

Other methods [7, 8] produce an *approximate envy-free* division, i.e., a division in each player feels their piece is within ϵ of being the best piece in their estimation. One such method, due to Simmons and described in [8], uses Sperner's lemma to accomplish a division by a minimal number of pieces. Hence it requires only $(n-1)$ cuts.

In spirit of work such as [5], we work with a discrete version of cake-cutting, which we call *necklace-cutting*. The only difference is that the set of points where one can make a cut is restricted to a discrete set, as when we cut a necklace of beads (opened at the clasp) we can only cut between the beads and not through beads, and the strings have no value so we can assume every string is only cut in the middle. We can model this situation by considering the cake as the interval $[0, l]$ where l is an integer and stating that only cuts at **integral** values are allowed, visualizing the beads at half-integers. In our problem, the obvious analogs of *divisions, cutset, prefers*, and *allocations* hold, so we do not redefine them. The *valuation* functions are now defined only on intervals of the form $[l_1, l_2]$ where $l_1 \leq l_2$ and both are integers. To remind ourselves of discreteness, we call the analogue of a piece of cake a *string of beads*. In terms of applications, it is natural to introduce a discrete analogue since many kinds of goods are indivisible, such as stamps, room assignments, or houses on a street. In terms of pure mathematics, it turns out that we have extra richness not found in the continuous case.

See Figure 2 for two ways of picturing the necklace. A recurring technique of this paper is to cut the "continuous cake analogue" of a necklace and then sliding the cuts to integral positions. Thus, we use the left representation for this remainder of this paper to stress the similarity between the two situations, although the right representation is more visually evocative of a necklace.

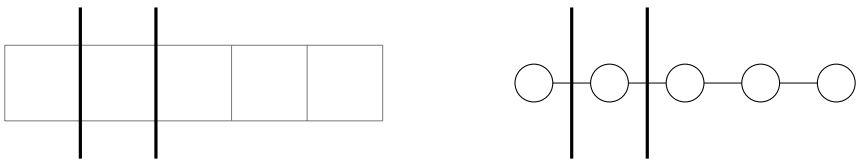

Figure 2: Two equivalent ways of picturing a discrete necklace with 2 cuts into 3 strings of beads. We use the left one from now on due to its similarity with the continuous case.

In the discrete case, we cannot extract power from the "continuity condition" as in the continuous situation. This should hint that our problem becomes harder. Indeed, it is no longer true that an envy-free cut must exist, even with a **arbitrary** number of cuts! The simplest counterexample is when we have a

single bead that everyone strictly prefers to having no beads. We cannot cut the bead, so one player must receive the bead and the others nothing. We can never achieve an envy-free cutting in this situation; furthermore, $(n-1)$ of the players are guaranteed to be envious in this case.

The lack of continuity means geometric and analytical techniques are more likely to fail. It is not a surprise, then, that the cake-cutting literature on discrete cakes is fairly sparse. For example, while Barbanel's encyclopedic variety of flavors of cake-cutting problems [2] is very careful about boundary conditions and alternate constraints for each problem, it already assumes by the end of the first page that the measures on the case do not contain point masses, thus ruling out the discrete case immediately. The Brams-Taylor text [4] occasionally mentions discrete items (using the terminology of *indivisible* goods) but generally stays within the realm of continuous cake-cutting.

Problem 1. *It seems natural to consider the following as fundamental problems for studying necklace-cutting:*

- *Unlike the continuous case, envy-free necklace-cutting is impossible in the most general case. What additional assumptions do we need to make to guarantee an envy-free division exists?*
- *What algorithm produces such divisions?*
- *When envy-free is impossible, what is the best we can do? For example, what is the smallest ϵ for which we can guarantee an ϵ-envy-free division?*

We examine these problems for increasingly general classes of constraints, starting with Section 2, where we give a proof of existence of an envy-free division in the case of *monolithic* preferences, slightly different from the one obtained by [5]. We then attack a more general version of the problem in Section 3. We end with a few glances into the future in Section 5.

2 Monolithic Preferences

As we frequently do in mathematics, we first study the simplest cases to get intuition. We start by studying a natural (but strong!) assumption, namely that each bead is only valuable to exactly one (or at most one, but beads that nobody gives positive valuations to can be trivially removed without changing the problem) player. If we are in such a situation, we say we have *monolithic preferences.* See Figure 3 for an example. In this situation, we can label each bead by the identity of the player, so we can denote a bead as an *P-bead* if only player P has a nonzero valuation of the bead. It turns out that we can indeed achieve an envy-free division in this case.

Theorem 2. */[5]/ Suppose we have a necklace with monolithic preferences. Then there exists an envy-free division for n players using only $(n-1)$ cuts.*

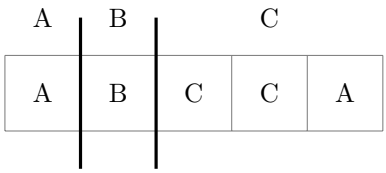

Figure 3: A discrete necklace with monolithic preferences, cut by 2 cuts into 3 strings of beads. The labels in the necklace denote the players who value those beads; the labels above the necklace denotes the allocation. We assume every player has valuation 1 for each bead for simplicity. This particular allocation is envy-free.

Proof. See [5]; the strategy uses Sperner's Lemma directly, simulating the proof of Theorem 1 as it is proven in [8]. □

This is a nice result to serve as the beginning of our investigations, and the rest of the paper is devoted to situations more general than the monolithic preferences case in two different directions: in Section 3 we generalize the valuation functions to have competing interests on every bead; in Section 4 we stay with monolithic preferences but generalize to cutting a 2-dimensional necklace. The common theme of the sections is that our proofs use the continuous result directly and then "shift" the cuts to integral positions. We do not know how to directly get our most general results with the Sperner's Lemma approach of [5] and [8].

3 General Preferences

While the mathematical result for the monolithic case is pleasant, the condition is fairly limiting; after all, it is usually the case that players bicker over things that more than one player wants! In this section, we explore what can be done in the case of *general preferences*, where any of the valuation functions $v_P(b)$ can be positive for any bead b.

Unfortunately, as we showed in Section 1, we cannot guarantee an envy-free division in all cases. Frequently players will bicker over a particular bead that cannot be further subdivided, such as in our 1-bead counterexample earlier, or in the slightly more complex example in Figure 4.

However, just because we cannot do envy-free does not mean we cannot come close. It makes sense to, as in [8], call a division ϵ-*envy-free* if for each player, the other pieces are worth no more than ϵ plus the piece assigned to him/her. As a special case, a division is 0-envy-free if and only if it is envy-free. In this case, we achieve the following result by using the continuous case as a tool, for **any** necklance:

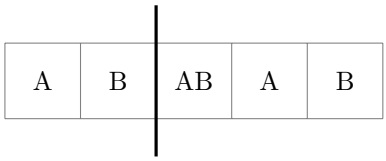

Figure 4: In this case of general preferences there is no envy-free division between 2 players; they will fight over the middle bead. The best we can do is a 1-envy-free division, cutting on either side of the middle bead.

Theorem 3. *Suppose we have general preferences. For a necklace of beads N where the value of each bead is at most s to every player, there always exists an ϵ-envy-free division of N among n players using only $(n-1)$ cuts, where $\epsilon < 2s$. In particular, there is always an $(2s)$-envy-free division.*

Proof. Represent our necklace N of k beads by a continuous cake N_C of length k, placed along the x-axis so that the leftmost point of the cake is at $x = 0$ and the rightmost point is at $x = k$. For each bead, make a length-1 cake where the value for each player who cares about the bead is uniform. For example, if we cut at $1/3$ from the left edge of the length-1 cake corresponding to bead B, then player P values the left piece $v_P(B)/3$ and right piece $2v_P(B)/3$, where $v_P(B)$ is the valuation of player P of bead B. There exists an envy-free division of N_C with $(n-1)$ cuts by Theorem 1; call it D.

In D, some cuts may be in middle of the cake pieces. Now, we round up or down each such cut to the nearest "integral" point of the cake. For example, if a cut is at $x = 1.2$, shift it so it is at $x = 1.0$. If a cut is at a half-integral length, then (important!) we always round to the right (or left; as long as we are consistent). Suppose player P were assigned a piece with value v in D. This process means P is now assigned a piece with value strictly greater than $v - s$ (he/she could have potentially had both ends of his/her assigned cake in D being made shorter, each end by at most $s/2$; this inequality is strict since the two ends cannot both contribute $s/2$ – since the cut to the right of the piece must have then been on a half-integral point and thus would have rounded right). Furthermore, any other piece in D could have increased, from the perspective of P, a value strictly less than s by the same logic. Thus, the envy for P after the process is strictly less than $2s$. □

Sometimes we are in situations where the valuations must be integral. In this case, the strictness of the inequality above can be quite useful. One such situation, which we will call *binary preferences*, is when each bead is worth exactly 1 or 0 to any player. In this case, we obtain:

Corollary 1. *Suppose we have binary preferences. For n players, there always exists an 1-envy-free division using only $(n-1)$ cuts.*

Proof. Theorem 3 shows that we can get the maximum envy to be strictly less than 2. However, since in binary preferences all valuations are integral (being

sums of 1's and 0's), the actual envy must be an integer. Thus, the envy is actually bounded above by 1. □

4 Two Dimensional Divisions, and a gerrymandering result

In this section, we consider a 2-dimensional grid of indivisible squares that can be cut along the edges[1]. The generality of Theorem 3 can be used to attack higher-dimensional continuous cake-cutting by projection-type arguments. We first explore what can be done with this method, then reach a stronger result in the monolithic preferences case.

An immediate corollary of Theorem 3 is:

Corollary 2. *Consider a 2-dimensional grid of squares with general preferences. Then there exists an ϵ-envy-free division among n players using $(n-1)$ parallel cuts, where $\epsilon < 2s$, s being the highest possible total valuation of the squares in any column by any player.*

Proof. Suppose our grid is $k \times l$. Project the onto an $1 \times l$ necklace by summing the valuations of the i-th column to obtain the valuation of the i-th bead. We are now in the situation of general preferences in 1 dimension and can apply Theorem 3. □

However, it seems that we can do significantly better if we allow our cuts to be non-vertical. In this section, define a *cut* to be a contiguous set of edges that separate the squares into two sets (as a Jordan curve) with its two ends on the top and bottom boundaries of the grid respectively. We consider two cuts to be *non-intersecting* if the cuts do not intersect each other transversely (i.e. there do not exist two points of a cut such that they are in the interiors of the two different sets created by the other cut).

Theorem 4. *Consider a 2-dimensional grid of squares with monolithic preferences. Then there exists an envy-free division among n players using $(n-1)$ non-intersecting cuts.*

Note that Theorem 2 is a special case of Theorem 4, by taking a $1 \times m$ grid for some m. Unfortunately, for larger grids, Theorem 4 does not promise that each piece is connected in an obvious way; two cuts can be non-intersecting even though they share some edges. It is possible to think of the pieces as "connected," but one must include the edges of its bordering cuts as part of the pieces. For an example, see Figure 5.

One might think of such a grid as a collection of indivisible districts in a state that we wish to divide among several constituencies. These constituencies

[1]Extending the equivalence from Figure 2, this model is of course equivalent with having indivisible beads at the center of the squares, with 4 strings coming out of each bead in the 4 cardinal directions orthogonally transverse to the edges of the squares.

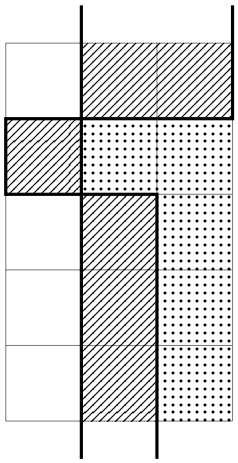

Figure 5: The grid is cut via 2 cuts into 3 pieces, each (barely) path-connected if we allow movement along the cuts themselves.

have some claim on the districts; perhaps some districts contain more of their constituencies. The net result is that we can produce somewhat gerrymandered regions that produce an envy-free division of the state.

Proof. We make the $(k \times l)$ grid into a continuous cake as we did in the proof of Theorem 3, by spreading the preference measure of each square uniformly in the square. Thus, there exist $(n-1)$ parallel vertical cuts that divide the grid cake into an envy-free allocation. We think of each of these vertical cuts as a union of k vertical edges. Assuming we have aligned our cake such that the lower-left corner is at the origin, call such a vertical edge *integral* if its x-coordinate is integral and *non-integral* otherwise. Our strategy is to slide non-integral vertical edges left or right so they become integral while keeping them connected via horizontal edges along the integral y-coordinates. We do this while holding the relative left-to-right order of the $(n-1)$ vertical edges on each horizontal strip constant, thus generating the desired cuts for the discrete grid of squares.

We say that a vertical edge *borders* players P and Q if the two adjacent pieces allocated by the original envy-free continuous allocation belong to P and Q. First, consider any square S desired by player P such that some vertical edge (possibly 2) going through S borders P. In this case, we can allocate the square completely to player P by moving the vertical edge (if it exists) bordering S on the left (and any other vertical edges going through S left of said edge) to the left of S and do the same for the right side of S. This gives P strictly more cake and does not affect the preferences and envy of other players, as only P cares about the square. Doing this for all squares ensures two conditions now hold for our (still envy-free) division D of the continuous cake: first, D assigns integral P-cake to every P (if not, then there must be some non-integral

vertical edge bordering P somewhere); second, if a square desired by P contains a non-integral vertical edge bordering P_1 and P_2, then neither P_1 nor P_2 can be P.

Our strategy is as follows. Suppose we have at least one non-integral vertical edge somewhere in our envy-free division. We describe a sliding process such that:

- we stay envy-free at all times,
- any time two non-integral vertical edges collide, we consider the two edges to have merged into a single edge (and sliding the resulting edge corresponds to the underlying edges moving together as a group), and
- any time a non-integral vertical edge becomes integral, we no longer slide them.

When two non-integral edges collide or a non-integral vertical edge becomes integral, the number of non-integral vertical edges decreases by 1; we can then repeat this process until all non-integral vertical edges become integral. Thus, it suffices to show that we can always perform this sliding process without losing the envy-free property.

Consider any square S with non-integral vertical edges going through it, desired by player P. Say that S is *contested* by Q if Q's piece contains part of S. Let the set of players contesting S contain P_1 and P_2 (recall that neither can be P). We say we *donate from P_1 to P_2 through S* if we slide the non-integral vertical edges through S in a (unique) way that the P_1 piece in S decreases at a constant speed, the P_2 piece in S increases at the same constant speed, and the other pieces in S stay at constant size. The key point is that since we are in an envy-free allocation, donating from P_1 to P_2 through S either causes the number of non-integral edges to decrease (which happens exactly when the P_1 part of S becomes empty), in which case we are done, or gets stuck when the P-cake allocated to P_2 equals the P-cake allocated to P, in which case any further donation would make P envious of P_2.

Suppose there exists at least one remaining square desired by player P with at least one non-integral vertical edge going through it. We define an auxiliary graph G with vertices indexed by players who are not P. For each square S desired by player P with at least one non-integral vertical edge, if S is contested by the player set T, draw an edge in G labeled by S for every pair of P_1 and P_2 such that P_1 and P_2 are both in T (we allow multiple edges in G). We have two cases:

- If we have a cycle (with no repeated vertices) in G of the form

$$P_1 \to P_2 \to \cdots \to P_k \to P_1,$$

  then we can simultaneously donate P_1 to P_2 (through the square corresponding to the edge between P_1 and P_2 in G), P_2 to P_3, etc. through P_k

to P_1. Since all players involved keep their P-cake amounts constant (having been donating and donated to at the same rate), at some point at least one of the parts belonging to some P_i in one of these squares becomes 0, corresponding to a decrease in the number of non-integral vertical edges.

- If we do not have a cycle, then some vertex must have degree 1. This means there is some square S desired by P and contested by $Q \neq P$ where Q does not have a fractional piece of a square desired by P anywhere else. This means the total amount of P-cake allocated to Q has fractional part exactly equal to the amount in S. Thus, we can give the entire square S to Q without fear that P will become envious of Q, as the P-cake allocated to P is currently integral.

In both cases, we can strictly decrease the number of non-integral vertical edges. Thus, we are able to perturb the vertical edges until all vertical edges are integral, in which case we have an envy-free allocation of the discrete grid. □

5 Conclusion

In this work, we attempted to attack heads-on the problem of fair-division when we remove the often-used constraint of "nice and continuous" measures. We borrowed ideas from continuous cake-cutting to attack discrete necklace-cutting, frequently thinking in terms of "shifting" a continuous division into a discrete one. It would be interesting to consider applications in reverse; for example, an ambitious (and possibly impossible) goal is leveraging these ideas to make progress toward a finite, bounded envy-free division of a continuous cake among n players.

Also, the idea of extending cake-cutting to the discrete case is very general; there is no need to be confined to the idea of envy-free divisions. There are many interesting questions we can ask. For example, Alon, Moshkovitz, and Safra [1] proved that if n players want to split a necklace with k kinds of beads such that each player gets as close to possible to the **the same number of beads**: $\lfloor a_i/n \rfloor$ or $\lceil a_i/n \rceil$ beads of the i-th kind, then they can do so using at most $(n-1)k$ cuts. This would be in the spirit of *proportional* divisions instead of envy-free divisions.

Two-dimensional cake/map cutting is relatively understudied in fair division. We hope Theorem 4 is a good step in this direction. The proof of Theorem 4 is short, but may have a more intuitive interpretation using graph theory, as the proof is reminiscent of network flow problems.

Finally, we mentioned in the introduction that the literature on discrete cake-cutting is relatively sparse, but that is only if we restrict ourselves to the particular framing of mathematical economics. It seems that much relevant mathematics to discrete cake-cutting can be found, albeit in different form, in other fields such as computer science and combinatorics where discreteness is the norm (such as the aforementioned work of Alon et al. [1]). Measure theorists and analysts with experience working with general measures probably also have

internalized intuition on this problem, just not with the same language. There is probably much untapped collaboration between these fields.

6 Acknowledgments

This work began at a 2014 AMS Mathematics Research Community, Algebraic and Geometric Methods in Applied Discrete Mathematics, which was supported by NSF DMS-1321794. Su was supported partially by NSF Grant DMS-1002938. We thank Boris Alexeev for valuable conversation.

References

[1] Noga Alon, Dana Moshkovitz, and Shmuel Safra, *Algorithmic construction of sets for k-restrictions*, ACM Transactions on Algorithms (TALG) **2** (2006), no. 2, 153–177.

[2] J.B. Barbanel and A.D. Taylor, *The geometry of efficient fair division*, Cambridge University Press, 2005.

[3] Steven J Brams and Alan D Taylor, *An envy-free cake division protocol*, American Mathematical Monthly (1995), 9–18.

[4] ———, *Fair division: From cake-cutting to dispute resolution*, Cambridge University Press, 1996.

[5] Javier Marenco and Tomás Tetzlaff, *Envy-free division of discrete cakes*, Electronic Notes in Discrete Mathematics **37** (2011), 231–236.

[6] J Neyman, *Un theoreme d'existence*, CR Acade. Sci. Paris **222** (1946), 843–845.

[7] J. Robertson and W. Webb, *Cake-cutting algorithms: Be fair if you can*, Ak Peters Series, Taylor & Francis, 1998.

[8] Francis Edward Su, *Rental harmony: Sperner's lemma in fair division*, American Mathematical Monthly (1999), 930–942.